

Is facial recognition technology worse at identifying darker-skinned faces than lighter ones?

RMIT ABC Fact Check
Updated 14 Feb 2020, 2:09pm

The claim

Facial recognition technology, which often relies on opaque algorithms developed by tech companies, is rapidly infiltrating many aspects of people's lives, despite lawyers and technology experts warning that it is inaccurate and discriminatory.

Human Rights Commissioner Edward Santow joined a chorus of voices critical of facial recognition technology, telling the ABC in an interview that it was prone to errors.

"Those errors are not evenly distributed across the community," he said.

"So, in particular, if you happen to have darker skin, that facial recognition technology is much, much less accurate."

He added: "When you use that technology in an area where the stakes are high, like in policing or law enforcement, the risks are very significant."

Mr Santow's comments came amid Morrison Government attempts to establish a centralised database for national security and official verification purposes that would include facial images drawn from documents such as passports and driver's licences, as well as other personal information.

Is Mr Santow correct in claiming the technology is less accurate at identifying darker-skinned faces than lighter ones? RMIT ABC Fact Check investigates.

The verdict

Mr Santow's claim is correct.

Research shows that facial recognition technology is less accurate when seeking to identify darker-skinned faces.

A study, which measured how the technology worked on people of different races and gender, found three leading software systems correctly identified white men 99 per cent of the time, but the darker the skin, the more often the technology failed.

Darker-skinned women were the most misidentified group, with an error rate of nearly 35 per cent in one test, according to the research, which was conducted by [Joy Buolamwini](#), a computer scientist at the Massachusetts Institute of Technology's (MIT) Media Lab.

Artificial intelligence (AI) systems such as facial recognition tools rely on machine-learning algorithms that are 'trained' on sample datasets.

So, if black people are underrepresented in benchmark datasets, then the facial recognition system will be less successful in identifying black faces.



PHOTO: [Facial recognition technology is better at identifying faces with lighter skin.](#) Ms Buolamwini's study states, "It has recently been shown that algorithms trained with biased data have resulted in algorithmic discrimination."

[Another study](#), which looked at the accuracy of six different face recognition algorithms found they were less accurate in identifying females, black people and 18-to-30-year-olds.

RMIT ABC Fact Check consulted a number of AI and legal experts who described Mr Santow's claim as correct, including internationally renowned expert [Toby Walsh, Scientia Professor of Artificial Intelligence at the University of NSW](#).

Professor Walsh, who is a Fellow of the Australian Academy of Science and a winner of the prestigious Humboldt research award, told Fact Check that despite more representative training data being used increasingly over the past 10 years, the technology was still less accurate in identifying darker-skinned faces.

"We're pretty certain it's not reliable at the minute," he said.

The context: what is the Australian government's plan for facial recognition?

The Morrison Government is seeking expanded use of facial recognition technology through its [Identity-matching Services Bill 2019](#) and the [Australian Passports Amendment \(Identity-matching Services\) Bill 2019](#).

It argues that the use of facial recognition technology would help ["prevent crime, support law enforcement, uphold national security, promote road safety, enhance community safety and improve service delivery."](#)

The bills would allow for the establishment of a centralised automated facial identification and verification system that would draw on a database holding photos from government-issued documents such as driver's licences, passports and citizenship data, as well as other personal information including names and addresses.

This would be known as the "interoperability hub". It would operate on a query and response model in which users of the system could submit a request for identity verification, but would not have access

to the underlying data powering the biometric processing. Users would simply receive a response to their query.

The underlying data would come under the control of the Department of Home Affairs which would maintain facilities for sharing facial images and other identity information between federal, state and territory government agencies, and, in some cases, private organisations, such as banks.

However, the bills faced a setback recently when the government-controlled Parliamentary Joint Committee on Intelligence and Security delivered a report which called for the Identity-matching Services Bill to be revamped to ensure stronger privacy protections and better safeguards against misuse.




PHOTO: [Is facial recognition technology prone to misuse?](#) (RMIT ABC Fact Check)

The spread of facial recognition systems

The global facial recognition market has grown rapidly over the past two decades and while estimates vary, [one market research company expects it to grow](#) in value from \$US3.2 billion (\$4.7 billion) a year to \$US7 billion (\$10.2 billion) in 2024.

In Australia, facial recognition technologies are used in airports as part of check-in, security and immigration procedures to speed up processing and reduce costs.

The technology is being used increasingly around the world:

- [The National Australia Bank and Microsoft](#) collaborated to trial a cardless ATM to enable customers to withdraw money using facial recognition and a PIN.
- [Westpac Bank](#) is exploring the use of AI technology to gauge the mood of its employees through their facial expressions.
- [Queensland Police](#) used facial recognition technology during the Gold Coast Commonwealth Games in 2018 but ran into considerable technical setbacks.
- [New Delhi police trialled facial recognition technology](#), scanning the faces of 45,000 children in various children's homes, and established and traced the identities of nearly 3000 children who had been registered as missing.
- [In China, police used](#) facial recognition technology to locate and arrest a man who was among a crowd of 60,000 concert goers.
- [The Chinese Government](#) is reportedly using facial recognition technology to track and control Uighurs, a largely Muslim minority.

'Arrest by algorithm'



[The China Cables leak of highly classified documents reveals the scale of Beijing's repressive control over](#)

Xinjiang, where more than a million Uyghurs and other Muslim minority groups are detained.

- This year, San Francisco became the first city in the US to ban police and other government agencies from using facial recognition technology to stop abuse of the technology amid privacy concerns.
- Amazon is marketing its facial recognition system, Amazon Rekognition, to police.

Despite the growing use of facial recognition technology, accuracy remains an issue. Police forces in London and South Wales have used the technology to identify and track suspects but both encountered high inaccuracy rates that were widely reported.

Liz Campbell, the Francine McNiff Professor of Criminal Jurisprudence at Monash University, told Fact Check the London Metropolitan Police and the South Wales force now used signage to indicate to the public when facial recognition technology was in use. "Even though, obviously, it caused quite a lot of consternation and turmoil, they're still being quite explicit when they're using it, whereas that's not the case in a lot of places in Australia," she said.

Professor Campbell said the technology was already in use in some supermarkets and petrol stations in Australia, where it was being used to monitor mood with the ultimate aim of targeting advertising to customers depending on their emotional state.

She said she had attended an AI conference in Melbourne in August where one of the speakers, Scott Wilson, the co-founder and chief executive officer of Wilson A.I., said facial recognition technology was used to detect emotion on the faces of delegates as they entered.

Schools in Australia also use facial recognition technology for security purposes, in which visitors' faces are photographed.

How does facial recognition technology work?

Biometrics is the application of statistical analysis to biological data. Facial recognition is a biometric technology capable of identifying or verifying a person by comparing the facial features in an image against a stored facial database.


The software first distinguishes a face from the background in a photo. It then examines the face closely and uses biometrics to map its features.

The technology measures the geometry of the face, using key markers such as the distance between a person's eyes, the shape of the nose or the length of the face to create a facial map, similar to a contour map of a landscape.

This unique facial signature, or 'faceprint', is then converted into data and checked against a database of stored faces in order to find a match.

A faceprint can be compared to an individual photo; for example, when verifying the identity of a known employee entering a secure office. Or, it can be compared to a database of many images in a bid to identify an unknown person.

Apart from facial recognition, AI-enabled recognition technologies include fingerprint, voice, iris and gait recognition systems which all have the potential to track and identify people. Facial recognition systems that identify emotions have also been developed, by companies such as [Affectiva/iMotions](#), [Microsoft](#), and [Amazon](#).



PlayGIF2.2 MBSettings

GIF: [Facial recognition technology has already been used by law enforcement in Australia.](#)

Lighter and darker skin — what the research shows

Global tech companies have produced commercial software that performs automated facial analysis used in many places, including schools, shopping centres, airports, local council areas, casinos and bars, and by law enforcement around the world.

Between 2014 and 2018, facial recognition software became 20 times better at matching a photo in a database of 26.6 million photos, according to a [report by the US National Institute of Standards and Technology \(NIST\)](#).

In an evaluation of 127 software algorithms from 45 different developers, the failure rate fell from 4 per to 0.2 per cent between 2014 and 2018, which the report said was largely due to "deep convolutional neural networks" that had driven an "industrial revolution" in face recognition.

But unlike the research conducted by MIT's Ms Buolamwini, the NIST study did not specifically measure accuracy in identifying photos of people with dark skin.

Ms Buolamwini, who campaigns to raise awareness of bias in algorithms, [conducted research](#) examining the performance of three leading facial recognition systems developed by Microsoft, IBM and Chinese technology company Megvii, to check how well they classified images by gender and race.

Her research involved building a database of 1270 photos of parliamentarians from six countries. She chose three African countries and three European countries to establish a range of skin colours from black to white and because of the gender parity in their national parliaments.

All three systems had relatively high accuracy rates, with Microsoft achieving the best result at 94 per accuracy overall.

Her research found that all three companies performed better on males than females, and also on lighter-skinned faces than darker-skinned ones. The systems performed worst on darker females with Microsoft recording an error rate of one in five (20.8 per cent) and

IBM and Megvii's Face++ more than one in three (34.7 per cent and 34.5 per cent respectively).

EMBED: [Error rates of facial recognition systems by skin colour and gender.](#)

IBM had the widest gap in accuracy, with a variance of 34 per cent in error rates between lighter males and darker females.

"I was surprised to see multiple commercial products failing on over one-in-three women of colour," [says Ms Buolamwini in a video about her research.](#)

She said more research needed to be done to find out exactly why the technology was less accurate with dark-skinned faces but, she added, "lack of diversity in training images and benchmark datasets" remained an issue.

[Another study](#), by five US computer scientists, analysed the accuracy of six different face recognition algorithms in relation to race/ethnicity, gender and age on a database of more than one million faces. It found they consistently performed worse on females, black people and 18-to-30-year-olds.

With regard to darker-skinned subjects, the study noted that while accuracy could be improved by using larger training datasets or datasets specific to a particular racial group, "even with balanced training, the black cohort still is more difficult to recognise."

It's worth noting that the study mentions there are well-documented factors that can affect the accuracy of facial recognition technology, including pose, illumination, expression, image quality and facial ageing.

It suggests the use of cosmetics by women may be one of the reasons why the algorithms were less accurate in identifying females.

Why does facial recognition technology have a racial bias?

While research shows facial recognition systems can have higher error rates when verifying darker-skinned faces, there is no consensus on the reason why.

There is, however, industry acknowledgement of racial bias in the algorithms as a result of unrepresentative training data.

[Dr Jake Goldenfein](#), a law and technology expert at Cornell Tech, Cornell University, USA, told Fact Check it was generally accepted that a lack of diversity in training data could affect the accuracy of facial recognition technology.

UNSW's Professor Walsh told Fact Check that despite improved training data, the technology was still not accurate when identifying darker-skinned faces.

The reasons, he said, were not fully understood, but he made two points: first, that humans tend to recognise the faces of people from their own race with greater accuracy than the faces of those from other races; and, second, black skin reflects light at different wavelengths to white skin.

"We don't fully understand [why], but we do know, for example ... white people do tend to recognise other white people better, whereas black people tend to recognise black people," he said.

"So, maybe it's very difficult to build a system that covers different races.

"From an optical perspective, black faces reflect light at very different wavelengths than white faces and trying to write one algorithm, I strongly suspect, is going to be, possibly, a fundamental challenge.

"It's not just the case that we need a more representative dataset. We've already got a more representative dataset and that didn't solve the problem."



PHOTO: [Could facial recognition technology be deployed widely in Australia?](#) (ABC News: Rebecca Trigger)

A recent [discussion paper on artificial intelligence commissioned by the Department of Industry Innovation and Science](#) and produced by CSIRO's Data61 team, makes a similar point.

"The challenge of ensuring fairness in algorithms is not limited to biased datasets," it says.

"The input data comes from the world, and the data collected from the world is not necessarily fair."

The paper says various sub-groups in society are represented differently across datasets, making AI "inherently discriminatory".

"Accuracy in datasets is rarely perfect and the varying levels of accuracy can, in and of themselves, produce unfair results — if an AI makes more mistake[s] for one racial group, as has been observed in facial recognition systems, that can constitute racial discrimination," it states.

Microsoft president Brad Smith, acknowledged in a [blog post](#) that some facial recognition technologies demonstrated "higher error

rates when determining the gender of women and people of color", saying work was under way in his company to improve accuracy. He called for governments to adopt laws to regulate the technology so that facial recognition services did not "spread in ways that exacerbate societal issues".

The Department of Industry, Innovation and Science this year drafted [voluntary ethics principles](#) to guide those who design, develop, integrate or use artificial intelligence, including a fairness principle which states "throughout their lifecycle, AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups".

How biased facial recognition systems can adversely affect black people

Professor Walsh told Fact Check the level of accuracy of facial recognition technology raised broader questions of fairness and accountability, particularly if it caused inconvenience for any racial group as a result of biased algorithms.

He said while it may seem like a bonus not to be recognised by facial recognition technology, in reality this could cause great inconvenience to people from non-white communities. He offered a number of examples where the technology's matching errors could affect — or were already affecting — black communities.

Some Australian schools, for example, were [using facial recognition technology to conduct the morning roll call](#) of students, although [Victorian state schools are banned from using facial recognition technology in classrooms](#) unless they have the approval of parents, students and the Department of Education.

Professor Walsh said that should that technology fail to identify Aboriginal students or other non-white students as being present, social services could be unnecessarily called in to investigate.

He said inaccuracies in identifying black people could lead to members of black communities being subjected to heavy-handed policing or security measures where their data was retained inappropriately.

In the UK, the Home Office — the department responsible for national security — [used face recognition technology for its online passport photo checking service despite knowing the technology failed to work well for people with dark skin](#), Professor Walsh pointed out.

"Again, that seems to be biased against people of colour, which is making their lives much more difficult," he said.

Darker-skinned people, who the system was unable to recognise, had been forced to attend an office to apply for a passport in person because they could not apply online like everyone else, he said.



PHOTO: [Facial recognition systems can make life more difficult for darker-skinned people.](#) (RMIT ABC Fact Check)

What is the legal framework around facial recognition technology?

Dr Sarah Moulds, a lecturer in law at the University of South Australia and an expert in human rights, counter-terrorism and anti-discrimination law, told Fact Check there was a very limited legal framework by which to protect people against misidentification by facial recognition technology.

"If the matching process is flawed or flawed more for [a specific racial] cohort, then the risk of a false positive increases and it is very difficult to challenge a false positive," she said.

"The existing data protection and human rights norms are applicable to face recognition technology, but the existing frameworks don't provide a sufficiently robust, nor sufficiently precise way of regulating face recognition technology," she said.

Principal researcher: [Sushi Das](#), chief of staff